IBM

Sales Guide
June 2006

IBM BladeCenter HS20

Product Overview

CONTENTS

Uncompromising 2-socket blade server for business-critical or telecommunications applications

[**Target Customers:** *Front-and mid-tier applications requiring high performance, enterprise-class availability and extreme flexibility.*]

Today's data center environment is tougher than ever. Customers are looking to reduce IT cost, complexity, space requirements, power consumption and heat output, while increasing flexibility, utilization and manageability. The **IBM® BladeCenter® HS20 blade server**, combined with the various BladeCenter chassis, can help them accomplish all of these goals.

A single **BladeCenter** or **BladeCenter H** chassis supports up to **14 hot-swappable** 30mm HS20 blades in only **7U** (BladeCenter) or **9U** (BladeCenter H) of rack space. (A rugged **8U BladeCenter T** chassis supports up to **8** HS20 blades in a **NEBS-3/ETSI-compliant** telecommunications environment.) In addition to the blade servers, these chassis also hold up to **four** (BladeCenter/BladeCenter T) **communication switches**, or up to **ten switches/bridges** (BladeCenter H) internally. Not only can this save significant data center space (and therefore the cost of floor space and rack hardware) compared to 1U servers, it also consolidates switches/bridges and cables for reduced complexity and lower cabling costs, and allows clients to manage everything in the solution as one. Using a BladeCenter chassis, up to **84** HS20 servers (**168** processors) can be installed in one **industry-standard 42U** rack but the value of BladeCenter extends far beyond high density data center environments.

Reducing an entire server into as little as **.5U** of rack space does *not* mean trading features and capabilities for smaller size. Each HS20 blade server offers features comparable to many 1U rack-optimized full-featured servers: Up to **two** of the latest **single-** and **dual-core** Intel® Xeon® processors, including full-performance/low-power models. The **single-core** Xeon processor is designed with **1MB** or **2MB** (model-specific) of **L2** cache, a leading-edge **800MHz** front-side bus (FSB), **64-bit extensions** (**EM64T**) and **Hyper-Threading Technology**, to help provide customers with the computing power they require to match their business needs and growth. The **dual-core** Xeon processor provides two cores with matching **1MB** L2 caches, for increased performance at a lower power draw than the single-core version. Both versions of the HS20 support up to **16GB** of **400MHz PC2-3200 ECC** (Error Checking and Correcting) double data rate II (**DDR2**) memory (in **4 DIMM slots**) with optional **Chipkill™** protection, for high performance and reliability. Integrated **dual Gigabit Ethernet**[1] controllers are standard, providing high-speed data transfers and offering **load-balancing** and **failover** capabilities. Via optional expansion cards, each blade can connect to additional Ethernet, **Myrinet**, **Fibre Channel**, **iSCSI**, **InfiniBand**™, and other high-speed communication switches housed in the chassis. Optional **2-port Expansion Cards** add additional fabrics to the HS20 server as needed.

All models offer impressive features at an equally impressive price, including up to **two Ultra320 SCSI** or **SAS** hard disk drives (model dependent, with **RAID-1** disk mirroring support. An optional 30mm **SCSI Expansion unit** connects to a blade (model-dependent) to provide an **additional two** SCSI HDDs, with **hot-swap** and enhanced **RAID-1E** support. Moreover, the HS20 is **optimized** for diskless operation, offering each blade server access to essentially unlimited storage capacity via Fibre Channel or iSCSI.

The various BladeCenter chassis are designed to monitor environmental conditions in the chassis and each blade and send alerts to the administrator. Advanced standard features, such as **Predictive Failure Analysis™,** light path diagnostics, **hot-swap redundant power supplies and blower modules** with **Calibrated Vectored Cooling™**; **IPMI 1.5** support, including **highly secure remote power on/off**; **text-console redirect over LAN**, a **management module** (upgradeable to hot-swap/redundant), **IBM Director** management software and **IBM**

---

[1] Actual data transfer speed will vary and is often less than the maximum possible. Gigabit Ethernet transfer speed requires support on both system and server, and appropriate network infrastructure.

1.

ServerGuide™ help maintain system availability with increased uptime.

If your customer needs highly manageable, high-performance computing power in a space-constrained environment, the HS20 is the ideal system.

Selling Features	**Price/Performance**

- The HS20 offers a choice of high-performance **Xeon** processors, from a **single-core** chip with 64-bit extensions, an **800MHz** front side bus and **1MB** or **2MB** of integrated Level 2 cache to a 32-bit **dual-core ultra-low power** Xeon processor with **double** the L2 cache **(2x1MB)** and an **800MHz** front side bus. This means there is an HS20 model to fit all budgets and power requirements.

- Selected HS20 blade servers are **NEBS3/ETSI-compliant** and feature long-life availability. These blades are ideal for telecom or Next Generation Network (NGN) applications such as IPTV, IP Multimedia Subsystem (IMS) and security.

- The **extremely high degree of integration** in the various BladeCenter chassis eliminates the need for internal server components, replacing numerous fans, KVM and Ethernet cables, power supplies, external switches and other components, with a few hot-swap/redundant components in the BladeCenter chassis itself. This integration also can greatly **reduce the amount of power consumed and heat produced**, relative to an equivalent number of 1U servers. This can significantly reduce a customer's power bill. The **reduced datacenter footprint** can also save on infrastructure cost.

- The chassis midplane provides **high-speed blade-to-blade, blade-to-switch-module and module-to-module communication** internally as well as externally. The midplane in the BladeCenter H provides **four 10Gb** data channels to each blade, and supports **4X InfiniBand** and **10Gb Ethernet** bridge modules.

- The BladeCenter chassis use **ultrahigh efficiency power supplies**. Most industry-standard servers use power supplies that are between **65-75% efficient** at converting power from AC wall current to the DC power used inside servers. BladeCenter power modules can be **more than 90% efficient**. This helps save even more money for the customer, because more of the power input is used for processing, rather than released into the data center as waste heat.

- BladeCenter also **reduces the number of parts required** to run the system. Sharing fans, systems management, floppy devices and media means fewer parts to buy, maintain, and fewer items that can bring the solution down.

Flexibility

The **HS20** has the ability to grow with a customer's application requirements, thanks to:

- An optional second Xeon processor.

- Up to **16GB** of high-speed **DDR2** error checking and correcting (ECC) system memory with optional Chipkill protection.

- Up to **two** internal **2.5**-inch SCSI HDDs, up to **two 3.5**-inch SCSI HDDs in an adjacent **Storage Expansion Unit**, and access to terabytes of external **IBM System Storage**™ and IBM **TotalStorage**™ SAN and NAS storage devices.

In addition, the various **BladeCenter chassis** offer a high degree of flexibility as well:

- Any combination of Xeon processor-based **HS20** and **HS40** (4-processor), blades can be used in the same chassis as Opteron processor-based **LS20** blades, and IBM PowerPC processor-based **JS20** and **JS21** blades. Up to **7 non-Low Voltage** HS20 blades or up to **8 Low Voltage** blades can be used in the BladeCenter T chassis. Depending on the blade servers used, the various BladeCenter chassis support Microsoft **Windows**, **Linux**, IBM **AIX**® and Sun **Solaris 10** operating systems in the same chassis.

- *Every LS20/HS20/HS40/JS20/JS21 blade server ever released by IBM is supported in every BladeCenter chassis ever release*d, going back to 2002. Every switch module released by IBM is the same way. (Ask HP and Dell how far back *their* compatibility goes.) Future blades and fabric switches are expected to continue to be compatible with previous chassis for the foreseeable future.

- A blade server has access to as many as **10 communication switches** and/or bridges in one **BladeCenter H** chassis. And the switches can be Ethernet, iSCSI, InfiniBand, Fibre Channel, Myrinet, or anything else designed and ServerProven for BladeCenter use. Switches, bridges and interface cards are currently available from such vendors as Brocade, Cisco, Intel, McData, Nortel, QLogic, Topspin and others, in addition to IBM.

Manageability

- The HS20 blade server includes a **Baseboard Management Controller (BMC)** to monitor

server availability, perform Predictive Failure Analysis, etc., and trigger IBM Director alerts.

- Each BladeCenter chassis includes a **Management Module** or **Advanced Management Module** to provide additional systems management capabilities, including *Web-based out-of-band control; virtual floppy and CD-ROM support; Windows "blue screen" error capture; LDAP and SSL support; and remote redirection of video, text, keyboard and mouse*.

- Integrated industry-standard **IPMI 1.5** support works with the BMC to alert IBM Director to anomalous environmental factors, such as voltage and thermal conditions. It also supports **highly secure remote power control**.

- IBM **Director 5.1x** is included for proactive systems management and works with both the blade's internal BMC and the chassis' management module. It comes with a portfolio of tools, including *Management Processor Assistant, RAID Manager, Update Assistant, Software Distribution and a Real Time Diagnostics tool*. In addition, IBM Director offers extended systems management tools for additional server management and increased availability. When a problem is encountered, IBM Director can issue administrator alerts via e-mail, pager, and other methods.

Availability and Serviceability

- BladeCenter chassis are designed for operation with **greatly reduced potential for single points of failure**. Most aspects of operation, from blade servers to communication modules, to management modules, to power and blower modules, are **hot-swappable**. The midplane connections are **redundant** and the other features can be made so, when used in pairs.

- HS20 blade servers support the use of optional **Chipkill-enabled memory DIMMs**. Chipkill memory can be up to **16X** better than standard ECC memory at correcting some types of memory errors. This can help reduce downtime caused by memory errors.

- **IPMI 1.5** supports highly secure remote system power on/off using data encryption. This allows an administrator to restart a server without having to visit it in person, saving travel time and getting the server back up and running quickly and securely.

- **Temperature-controlled blower modules** in the chassis adjust to compensate for changing thermal characteristics. At the lower speeds they draw less power and suffer less wear. Equally important in a crowded data center, temperature-controlled blowers produce less ambient noise in the data center than if they were constantly running at full speed.

- IBM **Real Time Diagnostics** allows the administrator to run diagnostics on a blade or the chassis without taking the hardware offline. This can provide greater server uptime.

- **Text console redirection** support allows the administrator to remotely view HS20 text messages over serial or LAN connections.

- A **standard three-year (parts and labor) limited onsite warranty**[2] gives your customers peace of mind and greater potential investment protection.

IBM Systems Agenda

The philosophy behind the design of the HS20 and the various BladeCenter Chassis involves four major aspects, Openness, Collaborative Innovation, Virtualization, and Technology Innovation that Matters:

- Unlike competitive blade solutions, many parts of the **BladeCenter specifications have been made open**. Hundreds of vendors have downloaded the specifications, and dozens so far have released products designed specifically for BladeCenter, including **Brocade**, **Topspin**, **Cisco**, **Nortel**, **Emulex**, and others. This provides customers with a wider selection of options to better fit their needs than a purely proprietary offering could.

- IBM has collaborated with industry leaders to offer **innovative solution packages for specific industry segments**, including banking, retail, telco and others. In addition, IBM formed **Blade.org**, a planned collaborative organization, to focus on accelerating the expansion of solutions for BladeCenter. Founding members in the blade.org community include Brocade, Citrix Systems, Intel Corporation, Network Appliances, Nortel, Novell and VMware.

- IBM offers virtualization capabilities in a number of ways: **desktop virtualization via Hosted Clients**, **InfiniBand virtualization** and **native blade virtualization** (using selected blades).

- Technology for technology's sake is pointless. Instead IBM BladeCenter focuses on technological innovation that provides real customer benefits. Examples of this include a **high degree of blade and module compatibility** across all chassis, the use of fewer and

---

[2] For terms and conditions or copies of the IBM Statement of Limited Warranty, call 800-772-2227 in the U.S. In Canada call 800-426-2255. Telephone support may be subject to additional charges. For warranties including onsite labor, a technician is sent after IBM attempts to resolve the problem remotely. International warranty service is available in any country in which this product is sold.

lower-power components to enable **significant power and cooling savings**, and hot-swap and redundant components that reduce **single points of failure**.

Key Features	**Intel Xeon Processors**

Intel Xeon Processors

The HS20 supports up to **two** identical Xeon processors (that is, the same clock rate and cache size and the same number of cores). The choice of processors includes:

- **2.8GHz single-core** Xeon processor with 64-bit extensions, an 800MHz FSB, **1MB** of L2 processor cache and **Hyper-Threading Technology**

- **2.8GHz Low Voltage single-core** Xeon processor with 64-bit extensions, an 800MHz FSB, **1MB** of L2 processor cache and **Hyper-Threading Technology** (NEBS 3/ETSI compliant)

- **3.0GHz Low Voltage single-core** Xeon processor with 64-bit extensions, an 800MHz FSB, **2MB** of L2 processor cache and **Hyper-Threading Technology** (NEBS 3/ETSI compliant)

- **2.8, 3.0, 3.2, 3.4, 3.6** or **3.8GHz single-core** Xeon processor with 64-bit extensions, an 800MHz FSB, **2MB** of L2 processor cache and **Hyper-Threading Technology**

- **1.67 or 2.0GHz ultra-low power** *dual-core* Xeon processor, with **667MHz** FSB and **2MB** of L2 processor cache (32-bit platform) (planned NEBS 3/ETSI compliant, 3Q/06)

Intel **Extended Memory 64 Technology (EM64T)** 64-bit extensions allow Xeon processors to use large memory addressing when running with a 64-bit operating system. This can result in higher performance. Additional registers and instructions (SSE3) can further boost performance for applications written to use them. Customers should contact their software provider to determine their software support for EM64T.

Some HS20 models feature new **dual-core Xeon** processors, which contain **two complete processor cores**, each with its own **1MB L2** cache. The two cores appear to software as two physical processors. The dual-core processors offer considerably higher performance than a same-speed Xeon processor with a single core.

Hyper-Threading Technology allows one physical processor to act as two logical processors that share resources, allowing two programs or threads to execute simultaneously and independently. By allowing the processor to use resources that would otherwise have been idle, Hyper-Threading Technology provides a performance boost of up to **30%** in CPU resource utilization (according to Intel) on multithreading and multitasking operations that can take advantage of the technology. Go to *http://www.intel.com/products/server/processors/server/xeon/index.htm* for more information on Hyper Threading Technology. The full effect of Hyper Threading Technology requires an operating system that has been optimized to take advantage of the feature. Go to *http://www.intel.com/support/platform/ht/os.htm* for a list of these operating systems.

DDR2 ECC Memory

The HS20 ships with **400MHz PC2-3200** double data rate II (DDR2) dual-interleaved memory DIMMs for fast access. It supports up to **16GB** of memory in **four** sockets, using industry-standard ECC (Error Checking and Correcting) registered DIMMs. It supports optional **Chipkill**[3] ECC technology (using 1GB or larger DIMMs), for **up to 16X** better error correction than standard ECC memory.

PC2-3200 memory is available in **512MB**, **1GB**, **2GB**, **4GB** and **8GB** memory **kits** (**two** 256MB, 512MB, 1GB, 2GB or 4GB DIMMs per kit, respectively). DIMMs are installed in pairs for increased performance provided by dual-interleaving.

Large HDD Storage Capacity

The HS20 offers a choice of disk storage, supporting up to **two 2.5-inch** internal **fixed Ultra320** drives.

- **10K** RPM **2.5-inch** — 36.4 or **73.4GB** capacities **(146.8GB maximum)**

If this is not enough internal storage, an optional "sidecar" storage blade is available. The **SCSI Storage Expansion Unit** is a 30mm blade that supports up to **two 3.5-inch hot-swap Ultra320 SCSI** HDDs. It is installed in the slot adjacent to an HS20 blade server.

- **10K** RPM **3.5-inch** — 73.4, 146.8 or **300GB** capacities **(600GB maximum)**

Fully populating both the HS20 blade and the expansion unit provides up to **746.8GB** of local storage to the HS20 blade server. If the customer needs more storage space, terabyte capacities are possible with optional external SAN storage solutions.
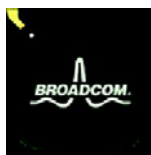
External Tape Storage

The HS20 supports various external rack-mounted SAN-attached tape drives. Supported tape technologies include:

- **IBM 2U Universal Tape Autoloader**
- **IBM 4U Universal Tape Library**

Gigabit Ethernet Controller

The HS20 includes an integrated **Broadcom 5704 dual-port** Gigabit Ethernet controller for up to 10X higher maximum throughput than a 10/100 Ethernet controller. The integrated controller offers **failover** and **load balancing** for better throughput and system availability. It also supports highly secure remote power management using IPMI 1.5, plus Wake on LAN® and PXE (Preboot Execution Environment) Flash interface.

Adapter Slots

The HS20 includes **one PCI-X** adapter slot on each blade. *Either* **one legacy BladeCenter expansion card** *or* **one SFF** (small form factor) BladeCenter adapter can be installed in this slot.

Adapters can be used to add interfaces to BladeCenter communication modules, including Fibre Channel, Myrinet, additional Gigabit Ethernet modules, InfiniBand, etc. Using an SFF adapter allows the use of both internal SCSI HDDs. Using a legacy card limits the number of usable internal HDDs to one.

The SCSI Storage Expansion Unit provides *two* additional adapter slots for SFF/legacy cards. (The first card must be a Gigabit Ethernet card.)

The *optional* **PCI I/O Expansion Unit II** provides *two* additional slots for standard PCI-X adapters.

---

[3] All models require optional Chipkill-enabled DIMMs for Chipkill protection. The DIMMs shipped with the systems are not Chipkill-enabled. Chipkill memory is designed to detect and correct single- and multibit errors. Reliability improvements depend on many factors and may be less than the maximum stated.

BladeCenter Chassis

IBM offers *three choices* of chassis in which to use the HS20 blade server. Each chassis serves different customer needs. The original **BladeCenter** chassis offers maximum density, great flexibility and a wide variety of expansion options at an entry-level price. The next-generation **BladeCenter H** chassis offers all of BladeCenter's capabilities, and adds new high-performance features. For those customers in need of a **ruggedized** chassis (for example, government/ military or telcos), **BladeCenter T** offers special features optimized for those environments. There is a high degree of interchangeability and compatibility of features among the three chassis. Any or all of these chassis can be installed in a rack along with other rack-optimized equipment.

BladeCenter is a **7U** chassis that provides:

- **Reduced single points of failure** — Many major components (either standard or optionally) are **hot-swappable** and/or **redundant**. Servers and modules can be configured for **automatic failover** to backups.

- **Forward and backward compatibility** — *Every blade (but one) and every switch and passthru module released by IBM for BladeCenter since 2002 is supported. New blades and fabric switches released by IBM for BladeCenter H or BladeCenter T will also be supported in the original BladeCenter chassis for the foreseeable future.*

- **Fourteen 30mm blade slots** — These **hot-swap** slots are capable of supporting up to **14** **HS20** or **HS21** (Xeon), **LS20** (AMD® Opteron®), and **JS20/JS21** (IBM PowerPC 970FX/MP) blade servers, or **7** double-wide (**60mm**) **HS40** blade servers or a **mixture** of 30mm and 60mm blades. It also supports 30mm optional **Storage and I/O Expansion Units, Memory and I/O Expansion Units** (for HS21), and/or **PCI I/O Expansion Unit IIs** in combination with the blade servers, using the same blade slots. Up to **six** chassis can be installed in an industry-standard **42U** rack, for a total of up to **84** 30mm blade servers per rack, *with full power redundancy*.

- **Four module slots** for **communication and I/O switches** — The modules interface with all of the blade servers in the chassis and eliminate the need for external switches or expensive, cumbersome cabling. All connections are done internally via the midplane. **Two** module slots are reserved for **hot-swap/redundant** Gigabit Ethernet switch modules. The other **two** bays support additional Gigabit Ethernet modules, or **Fibre Channel**, **InfiniBand,** and other switch modules or pass-through devices. All modules, when installed in pairs, offer **load balancing** and **failover** support. *Integrated* switch modules mean that **no extra rack "U space"** is required.

- **Two module bays** for **management modules** (IBM BladeCenter Management Module or the new Advanced Management Module, depending on the model of the BladeCenter chassis) — The management module provides advanced **systems management and KVM** capabilities for not only the chassis itself, but for all of the blades and other modules installed in the chassis. The Management Module provides capabilities similar to the IBM **Remote Supervisor Adapter II** used in stand-alone IBM System x™ or IBM eServer™ xSeries® rack and tower servers. The features of the Management Module can be accessed either **locally** or **remotely** across a network. One module comes standard. A second Management Module can be added for **hot-swap/redundancy** and **failover**.

- **Four module bays** for **Blower Modules** — Two **hot-swap/redundant** blower modules come standard with the chassis. They are capable of providing efficient cooling for up to **7 blades**. These modules replace the need for each blade to contain its own fans. The two blowers are more energy efficient than dozens or hundreds of smaller fans would be, and they offer many fewer points of potential failure. Two additional blower modules can be installed when more than 7 blades are used in a chassis.

- **Four module bays** for **Power Modules** — BladeCenter ships with two **2000W high-efficiency hot-swap/redundant** power modules (upgradeable to **four**), capable of handling the power needs of the entire chassis. Many servers use power supplies with an efficiency level of approximately 65-75%. Because BladeCenter uses power supplies that are at least **90% efficient**, much less power is wasted as heat. Not only is more power available for chassis use, there is less power wasted as excess heat output.

- **Redundant Midplane connections** — By giving each blade *two* physical connections to the midplane that connects all blades and modules together internally, a failure of one connector cannot bring down the server.

- **A hot-swappable Media Tray** containing a **DVD-ROM** drive, a **floppy drive**, **one USB 1.1** port, and a **light path diagnostic panel** — The media tray is *shared by all the blades* in the server. This eliminates unnecessary parts (and reduces the number of parts than can fail). In the event of a failure of the Media Tray the tray can be swapped for another. While the tray is offline, the servers in the chassis can remotely access the Media Tray in another chassis. The light path diagnostic panel contains LEDs that identify which internal components are in need of service.

It is extremely important to **include *all* infrastructure costs** when comparing a BladeCenter solution to a competitor's offering, not just the cost of the chassis and the blades. The high density and level of integration of the BladeCenter chassis greatly reduces the cost of the overall solution. For example, because up to **six** chassis will fit in a rack, this means that **up to 84 blade servers can be installed**. (Competitive blade offerings would require a second rack for this many blades.) Also, because up to **four** Ethernet, Myrinet, Fibre Channel, InfiniBand or other switches can be installed *per chassis*, **up to 24 switches can be installed per rack** *without* having to reserve any "U" space for switches, unlike the competition. (And the integrated switches may be less expensive than external, self-powered switches.) Plus, **the number of power distribution units (PDUs) needed per rack may be lessened**, because there are fewer discrete devices to have to plug in. In addition, because all the blades are connected to all the switches inside the chassis, **there is no need for external Ethernet or other communication cables to connect the blades and switches**. (Only the few cables needed to connect the switches to the external world are required.) This not only saves the cost of numerous cables per rack, but also the clutter and bother of routing that many cables. An added bonus is potentially **much freer airflow** behind the rack, due to fewer cables.

BladeCenter H is an even higher-performing **9U** chassis, designed for compute-intensive environments, such as Earth/Life Sciences, commercial analytics and next-generation network applications. Think of it as BladeCenter's big brother, with more speed and more features. It provides:

- **Reduced single points of failure** — Many major components (either standard or optionally) are **hot-swappable** and/or **redundant**. Servers and modules can be configured for **automatic failover** to backups.

- **Forward and backward compatibility** — *Every blade, switch and passthru module released by IBM for the original BladeCenter chassis since 2002 is supported in the BladeCenter H chassis.*

- **High-speed redundant Midplane connections** — Based on **4X InfiniBand**, the midplane supports up to **40Gb** bandwidth and provides **four 10Gb** data channels to each blade. By giving each blade *two physical connections* to the midplane that connects all blades and modules together internally, a failure of one connector alone cannot bring down the server.

- **Fourteen 30mm blade slots** — These **hot-swap** slots are capable of supporting any combination of **14 HS20** (Xeon), **LS20** (Opteron), and **JS20/JS21** (PowerPC 970FX/MP) blade servers, or **7** double-wide (**60mm**) **HS40** blade servers, or a **mixture** of 30mm and 60mm blades. It also supports 30mm optional **SCSI Storage Expansion Units** and/or **PCI I/O Expansion Unit IIs** in combination with the blade servers, using the same blade slots. Up to **four** chassis can be installed in an industry-standard **42U** rack, for a total of up to **56** 30mm blade servers per rack.

- **Up to ten module slots** for **communication and I/O switches or bridges** — The modules interface with all of the blade servers in the chassis and alleviate the need for external switches or expensive, cumbersome cabling. All connections are done internally via the midplane. **Two** module slots are reserved for **hot-swap/redundant** Gigabit Ethernet **switch** modules. **Two** slots support either **high-speed bridge modules** or legacy Gigabit Ethernet, **Myrinet**, **Fibre Channel**, **InfiniBand** and other **switch** modules. **Two** slots are dedicated for **bridge** modules. **Four** additional slots are dedicated for **high-speed bridge** modules. All modules, when in stalled in pairs, offer **load balancing** and **failover** support.

- *Integrated* switch and bridge modules mean that **no additional rack "U" space** is required.

- **Two module bays** for **Advanced Management Modules** — The new management module provides advanced **systems management and KVM** capabilities for not only the chassis itself, but for all of the blades and other modules installed in the chassis. The Advanced Management Module provides capabilities similar to the IBM **Remote Supervisor Adapter II** used in stand-alone xSeries rack and tower servers. New features include concurrent KVM (cKVM) and media tray, an external Serial over LAN connector, more memory, a more powerful onboard processor, industry-standard management interfaces (SMASH/CLP/CIM/HPI), USB virtualization, network failover and backward compatibility with the original Management Module, among others. The features of the module can be accessed either **locally** or **remotely** across a network. One module comes standard. A second module can be added for **hot-swap/redundancy** and **failover**. The module uses a **USB** connection, rather than the PS2 connection of the original Management Module.

- **Two module bays** for **Blower Modules** — Two **hot-swap/redundant** blower modules come standard with the chassis. They are capable of providing efficient cooling for up to **14 blades**. These modules replace the need for each blade and switch to contain its own fans. The blowers are more energy efficient than dozens or hundreds of smaller fans would be, and they offer many fewer points of potential failure. BladeCenter H also includes up to twelve additional hot-swap/redundant fans to cool the power supplies and high-speed switch

7.

modules.

- **Four module bays** for **Power Modules** — BladeCenter H ships with **two 2900W high-efficiency hot-swap/redundant** power modules (upgradeable to **four**), capable of handling the power needs of the entire chassis, including future higher-wattage processors. Each power module includes a customer-replaceable hot-swap/redundant fan pack (3 fans) for additional cooling capability.

- A **hot-swappable Media Tray** containing a **DVD-ROM** drive, **two USB 2.0** ports, and a **light path diagnostic panel** — The media tray is *shared by all the blades* in the server. This reduces unnecessary parts (and reduces the number of parts than can fail). In the event of a failure of the Media Tray the tray can be swapped for another. While the tray is offline, the servers in the chassis can remotely access the Media Tray in another chassis. The light path diagnostic panel contains LEDs that identify which internal components are in need of service.

- A **serial breakout port with optional cable** — This provides a direct serial connection to each blade server installed in the chassis, as an alternative to Serial over LAN. (***Note:*** This applies only to newer blades that include this capability.)

- It is extremely important to **include *all* infrastructure costs** when comparing a BladeCenter H solution to a competitor's offering, not just the cost of the chassis and the blades. The high density and level of integration of the BladeCenter H chassis can greatly reduce the cost of the overall solution. For example, because up to **four** chassis will fit in a rack, this means that **up to 56 blade servers can be installed**. Also, because up to **ten** Ethernet, Myrinet, Fibre Channel, InfiniBand or other bridges and switches can be installed *per chassis*, **up to 40 switches and bridges can be installed per rack** *without* having to reserve any "U" space for switches, unlike the competition. (And the integrated switches may be less expensive than external, self-powered switches.) Plus, **the number of power distribution units (PDUs) needed per rack may be lessened**, because there are fewer discrete devices to have to plug in. In addition, because all the blades are connected to all the switches inside the chassis, **there is no need for external Ethernet or other communication cables to connect the blades, bridges and switches**. (Only the few cables needed to connect the switches to the external world are required.) This not only saves the cost of numerous cables per rack, but also the clutter and bother of routing that many cables. An added bonus is potentially **much freer airflow** behind the rack, due to fewer cables.

BladeCenter T is a carrier grade, rugged **8U** (**20-inch** deep) chassis designed for challenging central office and networking environments. It provides:

- **NEBS 3/ETSI-compliance** — Designed for the Network Equipment Provider (NEP)/Service Provider (SP) environment. Also ideal for government/military, aerospace, industrial automation/robotics, medical imaging and finance.

- **Designed for Carrier-Grade Linux** — Several distributions are supported, include SUSE and Red Hat.

- **Reduced single points of failure** — Many major components (either standard or optionally) are **hot-swappable** and/or **redundant**. Servers and modules can be configured for **automatic failover** to backups. It also offers an extended product lifecycle (**3** years in production from date of General Availability, plus another **5** years of support).

- **Backward compatibility** — *Every blade, switch and passthru module released by IBM for the original BladeCenter chassis since 2002 is supported in the BladeCenter T chassis.*

- **Eight 30mm blade slots** — These **hot-swap** slots are capable of supporting any combination of **8 Low Voltage HS20** (Xeon) blade servers, or **7** regular-voltage **HS20**, **LS20** (Opteron), and **JS20/JS21** (PowerPC 970FX/MP) blade servers, or **4** double-wide (**60mm**) **HS40** or **Cell BE** processor-based blade servers, or a **mixture** of 30mm and 60mm blades. It also supports 30mm optional **SCSI Storage Expansion Units** and/or **PCI I/O Expansion Unit IIs** in combination with the blade servers, using the same blade slots. Up to **five** chassis can be installed in an industry-standard **42U** rack (or a telco rack), for a total of up to **40** 30mm blade servers per rack.

- **Four module bays** for **communication and I/O switches** — The modules interface with all of the blade servers in the chassis and eliminate the need for external switches or expensive, cumbersome cabling. All connections are done internally via the midplane. **Two** bays are reserved for **hot-swap/redundant** Gigabit Ethernet switch modules. The other **two** bays support additional Gigabit Ethernet modules, or **Fibre Channel**, **InfiniBand** and other switch modules. All modules, when in stalled in pairs, offer **load balancing** and **failover** support. *Integrated* switch modules mean that **no extra rack "U space"** is required.

- **Two module bays** for **Management Modules** — The management module provides advanced **systems management and KVM** capabilities for not only the chassis itself, but for all of the blades and other modules installed in the chassis. The Management Module provides capabilities similar to the IBM **Remote Supervisor Adapter II** used in stand-alone

8.

xSeries rack and tower servers. The features of the Management Module can be accessed either **locally** or **remotely** across a network. One module comes standard. A second Management Module can be added for **hot-swap/redundancy** and **failover**. The modules include a light path diagnostics panel containing LEDs that identify which internal components are in need of service.

- **Four module bays** for **Blower Modules** — **All four hot-swap/redundant** blower modules come standard with the chassis. These modules replace the need for each blade to contain its own fans. The blowers are more energy efficient than dozens or hundreds of smaller fans would be, and they offer many fewer points of potential failure.

- **Four module bays** for **Power Modules** — BladeCenter T ships with **two 1300W high-efficiency hot-swap/redundant DC** or **AC** (model-specific) power modules (upgradeable to **four**), capable of handling the power needs of the entire chassis.

- **Redundant Midplane connections** — By giving each blade *two* physical connections to the midplane that connects all blades and modules together internally, a failure of one connector alone cannot bring down the server.

- **A Tray** containing a **CD-ROM** drive, **KVM ports**, **two USB 1.1** ports, and a **light path diagnostic panel** — The media tray is *shared by all the blades* in the server. This eliminates unnecessary parts (and reduces the number of parts than can fail). A **floppy drive** is optional.

- **Long-life availability** — The BladeCenter T chassis will be produced for at least 3 years from date of General Availability. This allows telecom Network Equipment Manufacturers (NEPs) and, Service Providers (SPs) to standardize on a configuration for longer than traditional enterprise platforms. Product availability for periods longer than 3 years will be handled on an individual basis.

- It is extremely important to **include *all* infrastructure costs** when comparing a BladeCenter T solution to a competitor's offering, not just the cost of the chassis and the blades. The high density and level of integration of the BladeCenter T chassis can greatly reduce the cost of the overall solution. For example, because up to **five** chassis will fit in a rack, this means that **up to 40 blade servers can be installed**. Also, because up to **four** Ethernet, Fibre Channel or other supported switches can be installed *per chassis*, **up to 20 switches can be installed per rack** *without* having to reserve any "U" space for switches. (And the integrated switches may be less expensive than external, self-powered switches.) Plus, **the number of power distribution units (PDUs) needed per rack may be lessened**, because there are fewer discrete devices to have to plug in. In addition, because all the blades are connected to all the switches inside the chassis, **there is no need for external Ethernet or other communication cables to connect the blades and switches**. (Only the few cables needed to connect the switches to the external world are required.) This not only can save the cost of numerous cables per rack, but also the clutter and bother of routing that many cables. An added bonus is potentially **much freer airflow** behind the rack, due to fewer cables.

Light Path Diagnostics

Light path diagnostics enables a technician to quickly identify and locate a failed or failing system component, such as a specific blower module or memory DIMM. This enables quick replacement of the component, which helps increase server uptime and lower servicing costs.

The front of each blade server—and the chassis itself—has an LED indicator light to show possible component failures. This lets the servicer identify the failing component without the need to or remove the blade server from the chassis. The light path diagnostics panel tells the servicer which component of the affected server requires attention.

In addition, many components have their own identifying LEDs. For example, each of the memory modules has an LED next to the socket, as do both processors. This allows the servicer to easily identify exactly which component needs servicing. By following the "light path," the component can be replaced quickly, and without guesswork. (***Note:*** In the event of a failed DIMM, the system will restart and mark the DIMM as bad while offline, thus allowing the system to continue running, with reduced memory capacity, until serviced.)

Advanced Systems Management Capabilities

Each BladeCenter chassis offers a high level of systems management capabilities that are well-suited to remote locations as well as to stand-alone environments. Features include the Management Module (MM) or Advanced Management Module (aMM), Baseboard Management Controller (BMC), Automatic Server Restart, PowerExecutive™, Wake on LAN® support, PXE 2.0 support, text and graphics console redirect, Real Time Diagnostics, Predictive Failure Analysis, IBM Director and Remote Deployment Manager.

The **MM/aMM**, in combination with the HS20 blade server **BMC**, provides industry-standard **Intelligent Platform Management Interface** (**IPMI**) **1.5**-compliant systems management. It

provides a number of important system functions, including:

- Monitoring of system and battery voltage, system temperature, fans, power supplies, processor and DIMM status
- Fan speed control
- Product ID and Family ID detection
- Highly secure remote power on/off
- System reset control
- NMI/SMI detection and generation
- System diagnostic LED control (power, HDD, activity, alerts, heartbeat)
- IPMI over LAN
- Proxy server support
- LAN messaging and alerting
- Local update of BMC firmware
- Support for IPMI v1.5 compliant management software (e.g., xCAT)
- Other mandatory and optional IPMI BMC functions

The BMC, via the management module, alerts IBM Director to anomalous environmental factors, such as voltage and thermal conditions—even if the server has failed.

Other systems management features offered for the combination of blade server and chassis include:

- Predictive Failure Analysis for system processors, memory and HDDs, as well as chassis switch modules, blower modules and power modules
- Web-based out-of-band control
- Windows "blue screen" capture
- Remote virtual media
- High-speed remote redirection of PCI video, keyboard and mouse
- SSL (Secure Socket Layer) and LDAP (Lightweight Directory Access Protocol) support

In order to put control of processor power-saving features at the fingertips of administrators, IBM developed **PowerExecutive**. PowerExecutive is designed to take advantage of new processor features, such as balancing the performance of the system according to available power input. It provides the ability to plan and predict power consumption based on your BladeCenter hardware configuration. It also allows you to reduce the infrastructure required for redundancy, by using fewer servers on smaller power feeds and potentially lowering your overall data center support costs. It does this by inventorying all components at the blade level, then adding up the power draw for each blade and tracking that usage. In failure mode, PowerExecutive (through the BladeCenter Management Module) might request that certain blades in each domain throttle down to reduce power consumption.

Automatic Server Restart (ASR) helps reduce downtime by restarting the server automatically in the event of a system lockup. ASR technology is a combination of hardware circuitry tied into the server's system reset function and a device driver. As long as the server continues running, the ASR watchdog timer will keep being reset, but if the operating system crashes or the hardware freezes somehow the ASR software will be unable to reset the hardware timer. If the timer is not reset within five minutes, it automatically triggers the ASR hardware, which immediately restarts the server (and logs an ASR event with IBM Director). These features are designed so that *no more than five minutes can pass before the server is restarted*.

Text and Graphics Console Redirect support allows the administrator to remotely view HS20 text and graphics messages over serial or LAN.

Wake on LAN permits the HS20 to be remotely powered on if it has been shut off. Once powered up, the server can be controlled across the network, using the **Preboot Execution Environment** (PXE).

Like Wake on LAN, **PXE** is system firmware. It allows software such as the **IBM Remote Deployment Manager** to take control of a system before the BIOS, operating system or applications are loaded (using Wake on LAN/PXE) and lets an administrator perform many low-level tasks remotely that would otherwise require a visit to each system. These tasks may include such things as formatting a hard disk drive, updating system firmware, or deploying a Windows or Linux operating system.

Real Time Diagnostics allows an administrator to perform remote diagnostics on any blade server or any chassis without downtime.

Predictive Failure Analysis (PFA) enables the MM/aMM and the BMC to detect impending

10.

failure of supported components (processors; memory; expansion cards; switch, blower and power supplies; and hard disk drives) as much as 48 hours before actual failure, and alert the administrator through IBM Director. This gives customers the ability to replace the failing component *before* it fails, resulting in increased uptime.

IBM Director 5.1x software for advanced workgroup management is included with the HS20. IBM Director comes with a portfolio of tools, including Management Processor Assistant, Rack Manager, RAID Manager, Update Assistant and Software Distribution. System Availability (a no-charge download) and Capacity Manager (sold separately) are available as add-ons for additional server management and increased availability. IBM Director provides a single uniform graphical interface for all of these systems management functions.

IBM Director enables the customer to customize thresholds and monitor system components (for things like temperature, voltage regulation, etc.) to help maximize uptime.

Extensive System Support Features

The IBM services and technical support portfolio provides world-class, consistent, high-quality service and support. From the start, IBM programs make it easier for customers to plan for, configure and purchase BladeCenter servers, get them running and keep them running long-term. These features include IBM Express Portfolio, IBM ServerProven®, the IBM xSeries and BladeCenter Rack Configuration Tool, IBM ServerGuide, IBM Electronic Service Agent™, Product Customization Services and extensive technical support offerings.

Select configurations of xSeries servers are part of the **IBM Express Portfolio**, designed, developed and priced to meet the specific needs of midsized businesses. The IBM Express Portfolio of solutions are easy to acquire, install and manage. And they leverage IBM technology to provide tangible solutions to help you solve business problems in an on demand world.

The IBM **ServerProven** program provides the confidence that specific options and operating systems have been tested on the blade servers and are officially supported to work together. It is updated frequently to keep the latest compatibility information at your customers' fingertips.

IBM **xSeries and BladeCenter Rack Configurator** is a downloadable tool that simplifies the often complex chore of configuring a full rack of servers and confirming that customers have all the cables, power distribution units, KVM (keyboard, video and mouse) switch boxes and other components they need, as well as the proper airflow clearances, electrical circuits and other environmental conditions.

IBM **ServerGuide** (installed from CD) simplifies the process of installing and configuring xSeries servers. ServerGuide goes beyond mere hardware configuration by assisting with the automated installation of the Microsoft® Windows® Server 2000 and 2003 operating systems, device drivers and other system components, with minimal user intervention. (Drivers are also included for support of Novell NetWare, Red Hat Linux and SUSE LINUX.) This focus on deployment helps customers reduce both their total cost of ownership and the complexity that administrators and technical personnel face.

IBM offers an innovative "call home" feature that allows xSeries servers to automatically report hardware problems to IBM support, which can even dispatch onsite service[4] if necessary to those customers entitled to onsite support under the terms of their warranty or an IBM Maintenance Agreement. The IBM **Electronic Service Agent** is a downloadable software tool available from the IBM support Web site at no extra charge. It resides on a server and provides electronic support and problem management capabilities through a highly secure electronic dialogue between customer systems and IBM. Electronic Service Agent monitors networked servers for hardware errors and it can perform hardware and software inventories and report inventory changes to IBM. All information sent to IBM is stored in a highly secure database and used for improved problem determination.

Additional services include hardware warranty upgrades and factory-installed **Product Customization Services** (PCS), such as asset tagging, hardware integration, software imaging and operating systems personalization.

IBM offers extensive **technical support** by phone and via the Web. Support options include links to forums/newsgroups, problem submission, online shopping support, service offerings, device drivers for all IBM product lines, software downloads and even upcoming technical seminar worldwide schedules and registration. Also available are remote installation, configuration and usage support for both xSeries hardware and software, as well as onsite custom services to give customers the level of expertise they require.

---

[4] For onsite labor, IBM will attempt to diagnose and resolve the problem remotely before sending a technician.

Please see the Legal Information section for important notices and information.

Key Options	**IBM options for xSeries servers let customers take their servers to a higher level**

Make sure that customers know they can rely on xSeries and blade options to supply a comprehensive solution for their business needs. Options help them create an optimized server system to meet their data protection, storage and availability needs. Every IBM option is designed and tested for peak performance and flexibility, helping to maximize the customer's return on investment. The combination of xSeries servers and options lets customers keep their fingers on the pulse of their e-business.

Processors — The Intel Xeon processor provides high clock rates, 64-bit extensions, a large cache, Hyper Threading Technology and advanced features for availability and manageability. Large cache size, combined with an **800MHz** system bus speed, reduce memory latency and facilitate the movement of data through the processor and I/O devices. (**Note:** System performance depends not only on the number of processors in the server but also on the power and functionality of each processor.) Let your customer know that adding a second processor may be a cost-effective way to achieve significant performance improvements.

Memory — Memory is a significant factor in systems application performance. Adding more memory to a BladeCenter server is one of the most effective ways to increase application performance. For best performance in a server with a dual-core processor, there should be twice as much memory available as for a single-core processor.

Hard Disk Drives — IBM hard disk drives help customers improve the transaction and cost performance of their HS20 servers. The choice of hard disk drives can be a critical aspect of maximizing the I/O throughput of the system. **Ultra320 SCSI** hard disk drives (**2.5-inch**) are available for the HS20 with capacities up to **73.4GB** at **10,000** RPMs. Additionally, a **SCSI Storage Expansion Unit** can be attached to the HS20 to add up to two **3.5-inch** Ultra320 SCSI drives of up to **300GB** apiece at **10,000** RPMs.

I/O slots — The HS20 supports the addition of a **PCI I/O Expansion Unit II**, which provides **two front-accessible standard PCI-X slots**. In addition, the **SCSI Storage Expansion Unit** provides **two SFF** (small form-factor) adapter slots.

External Storage — The **IBM TotalStorage DS300** offers remote SAN storage using the **iSCSI** storage protocol across IP networks. The DS300 scales up to **4.2TB** of storage capacity using standard hot-swap Ultra320 SCSI HDDs in a 14-drive 3U cabinet. Three models offer single or dual RAID controllers (with each controller supporting two Gigabit Ethernet ports), up to 256MB or 1GB of cache memory and dual redundant power supplies and fans. For low-cost **Fibre Channel SAN** storage, the **DS400**, when combined with two EXP400 storage enclosures, can scale up to **12.6TB** of hot-swap SCSI storage, incorporating redundant power, cooling and **dual 2Gbps** Fibre Channel controllers.

Midrange **TotalStorage DS4000** products are designed to support the large and growing data-storage requirements of business-critical applications. Because DS4000 solutions can support connectivity distances of up to **10km (6.2 miles)** at Fibre Channel rates of up to **2Gbps**, companies can more easily configure offsite xSeries servers and storage systems to keep critical data available around the clock—even in the event of a catastrophe.

These storage servers are designed to provide highly-available Fibre Channel RAID protected storage that provide the foundation for Storage Area Networks (SANs). **DS4500** storage servers deliver scalability up to **32.8TB** of Fibre Channel storage or **56TB** of SATA storage (**DS4500** only). They provide end-to-end Fibre Channel solutions with high-availability fault-tolerant components.

Additionally, external LAN-attached tape storage is available.
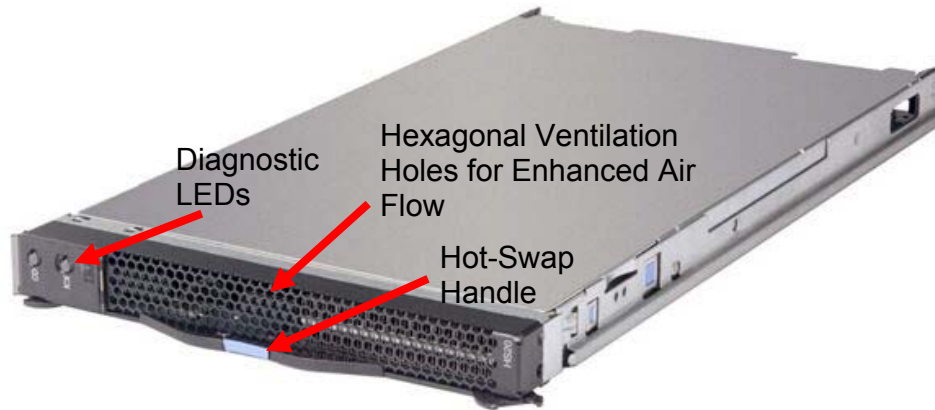
Communication Modules — The various BladeCenter chassis support integrated communication and I/O switches and/or bridges for Gigabit Ethernet, Myricom, Fibre Channel, InfiniBand, and others. Expansion adapters for individual HS20 blades are available to interface with these modules.

Rear Door Heat eXchanger — The unit attaches to the back of an IBM S2 42U Enterprise Rack. It is capable of removing up to 50,000 BTUs (14KVa) from the data center using water lines under the raised floor. The door swings open for servicing.
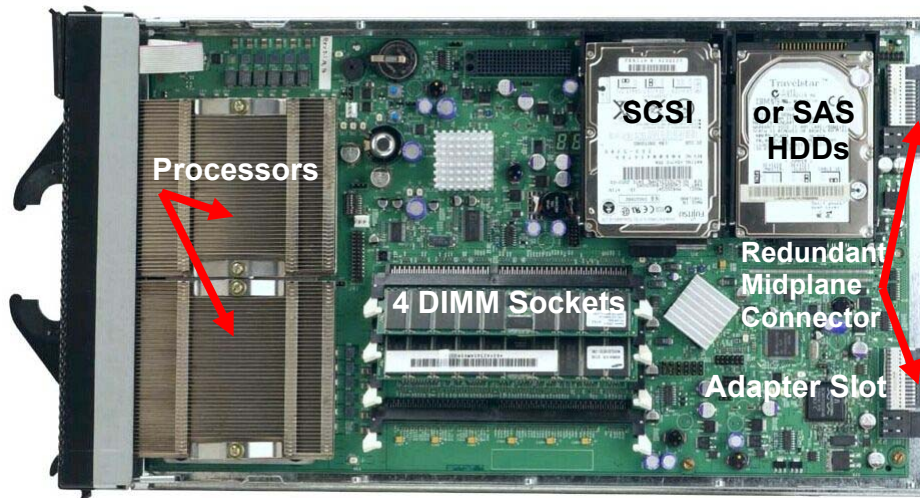
Redundant features — Optional power supply modules, blower modules, management modules, switches and bridges provide redundancy for the various BladeCenter chassis.

Workhorse 2-socket single- or dual-core Intel Xeon blade server

HS20 Images

Front View

Diagnostic LEDs

Hexagonal Ventilation Holes for Enhanced Air Flow

Hot-Swap Handle

Interior View

Processors

SCSI

or SAS HDDs

Redundant Midplane Connector

4 DIMM Sockets

Adapter Slot

BladeCenter HS20 Specifications			
Machine type	8843-0xX/0xY, 2xX/2xY, 3xX/3xY, 4xX/4xY, 5xX/5xY, ExX/ExY, JxX/JxY, LxX/LxY, MxX/MxY		
Form factor	30mm blade		
Processor type	**Single-core** Intel **Xeon** processor with **EM64T** **2.8GHz** (0xX/0xY, E7X/E7Y, LTX/LTY), **3.2GHz** (2xX/2xY, E2X/E2Y), **3.4GHz** (3xX/3xY) **3.6GHz** (4xX/4xY), **3.8GHz** (5xX/5xY, E9X/E9Y)	**Single-core Low-Voltage** Intel **Xeon** processor with **EM64T** **2.8GHz** (JxX/JxY, L1X/L1Y) , **3.0GHz** (MxX/MxY)	**Dual-core** Intel **Xeon** processor with **EM64T** **1.67GHz** (2FX/2FY**), 2.0GHz** (3FX/3FY)
Processor power draw	105W	**55W**	**31W**
Front-side bus (FSB) speed	800MHz		
# of processors standard / maximum	1 / 2		

13.

BladeCenter HS20 Specifications				
NEBS/ETSI Compliance	Select blades are NEBS3/ETSI-compliant. (Contact IBM sales for details.)			
Internal L2 cache	**1MB** (4TX/4TY, JxX/JxY, LxX/LxY)		**2MB** (all other models)	
Chipset	Intel **E7520**			
Standard / maximum memory[5]	**512MB** (2 x 256MB) / 16GB (x1X/x1Y, x7X/x7Y)	**1GB** (2 x 512MB) / 16GB (x5X/x5Y, xTX/xTY)	**2GB** (2 x 1GB) / 16GB (x2X/x2Y, 2FX/2FY, 3FX/3FY)	**4GB** (2 x 2GB) / 16GB (x9X/x9Y)
Standard memory type	PC2-3200 (**400MHz**) DDR II ECC			
Memory interleaving	Yes (using 2 or 4 DIMMs)			
DIMM capacities supported	256MB, 512MB, 1GB, 2GB, 4GB			
Chipkill protection supported	Yes (using 1GB or larger DIMMs)			
# of DIMM sockets total / available	4 / 2			
# of 2.5-inch drive bays total / available	**2 / 2 Ultra320 SCSI** (27X/27Y, E2X/E2Y)		**2 / 1 SAS** (all other models)	
# of 3.5-inch drive bays total / available	**2 / 2** (all other models, using optional SCSI Storage Expansion Unit)		**None** (2FX/2FY, 3FX/3FY)	
Maximum internal 2.5" HDD capacity	**146.8GB** (2 x 73.4GB) **fixed Ultra320 SCSI** (using SFF card or no expansion card); **73.4GB** (1 x 73.4GB) **fixed Ultra320 SCSI** (using legacy expansion card) — all other models		**146.8GB** (2 x 73.4GB) **fixed SAS** (using SFF card or no expansion card); **73.4GB** (1 x 73.4GB) **fixed SAS** (using legacy expansion card) — 2FX/2FY, 3FX/3FY	
Maximum 3.5" HDD capacity	**600GB** (2 x **300GB**) **hot-swap** Ultra320 SCSI (all other models)		**None** (2FX/2FY, 3FX/3FY)	
2.5-inch HDD capacities supported	36.4, **73.4GB** — **10K** RPMs			
3.5-inch HDD capacities supported	73.4, 146.8, **300GB** — **10K** RPMs		N/A	
# of HDDs standard	1 x **36.4GB** (27X/27Y, E2X/E2Y)		None (all other models)	
# of optical drives standard	None (one standard in chassis)			
# of diskette drives standard	None (one standard in BladeCenter/BladeCenter H chassis)			
Internal tape drives supported	None (SAN-attached)			
Disk drive technology	**Ultra320** SCSI (all other models)		**SAS** (2FX/2FY, 3FX/3FY)	
Integrated disk controller	Single-channel LSI Logic 53C1020 (all other models)		Single-channel LSI Logic 53C1066 (2FX/2FY, 3FX/3FY)	
# of disk drives supported *per channel*	4 (four drive limit, due to available bays)			
External disk drive support	NAS/SAN-attach			
# of adapter slots total / available	**1** BladeCenter-unique SPF PCI-X slot on the blade; **2** additional PCI-X slots available with the optional **PCI Expansion Unit II**			
# of 64-bit / PCI-E slots	None			
# of 64-bit / PCI-X slots	None (2 via optional PCI Expansion Unit II)			
# of 32-bit / PCI slots	None			
# of video ports	None (chassis-attached)			

---

[5] Maximum memory and disk capacity may require the replacement of standard components with the largest supported component available.
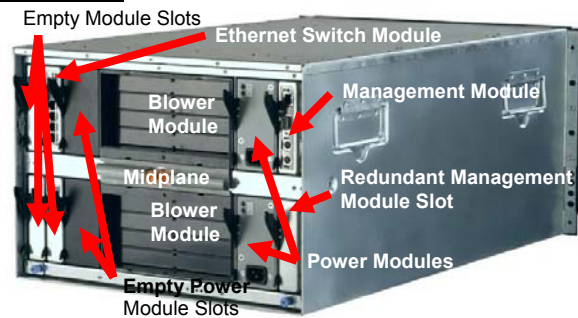
Please see the Legal Information section for important notices and information.

BladeCenter HS20 Specifications	
Video controller	ATI RADEON 7000M
Video memory	16MB SDRAM
Maximum video resolution at 32-bit color	1024 x 768 x 32-bit color at 85Hz
Gigabit Ethernet controller	Broadcom 5704
# of Gigabit Ethernet ports	2
# of RS485 ports	None
# of serial ports	None (1 direct via BladeCenter H chassis, or Serial over LAN in BladeCenter and BladeCenter H)
# of parallel ports	None
# of mouse ports	None (1 via chassis)
# of keyboard ports	None (1 via chassis)
# of USB ports	None (2 via chassis)
Systems management controller	Integrated BMC
Diagnostic LEDs (front panel)	Power good, blade location, over temperature, information, general fault
Predictive Failure Analysis support	Processor, memory, HDDs, expansion cards
Power supply size	Contained in chassis
# of power supplies standard / maximum	Contained in chassis
# of fans/blowers standard / maximum	Contained in chassis
Dimensions (HWD) / weight	9.7" (245mm) **H** 1.14" (29mm) **W** 17.6" (446mm) **D** 12 **lbs** (maximum) 5.44 **kg** (maximum)
Operating systems supported	Microsoft Windows Server/Advanced Server 2000, Microsoft Windows Server 2003 (Standard/Web/Enterprise Editions), Microsoft Windows XP (Client certification), RHEL 3/4, SLES 8/9, Novell NetWare 6.5, Sun Solaris 10
Extended/long life support	Select blades are supported for long life. (Contact IBM sales for details.)
Length of limited warranty	3 years (parts and labor) onsite

BladeCenter Chassis Images

Front View

CD-ROM Drive
Media Tray
Blade Server
Floppy Drive
Other Blade Servers
Light Path Diagnostics Panel

Rear View

Empty Module Slots
Ethernet Switch Module
Management Module
Blower Module
Redundant Management Module Slot
Midplane
Blower Module
Power Modules
Empty Power Module Slots

Interior View

Diagnostic LEDs   USB Port   CD-ROM Drive   Floppy Drive

Redundant Midplane
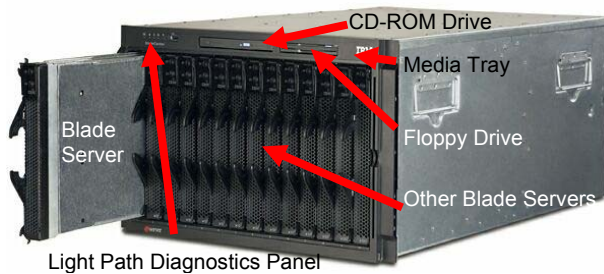
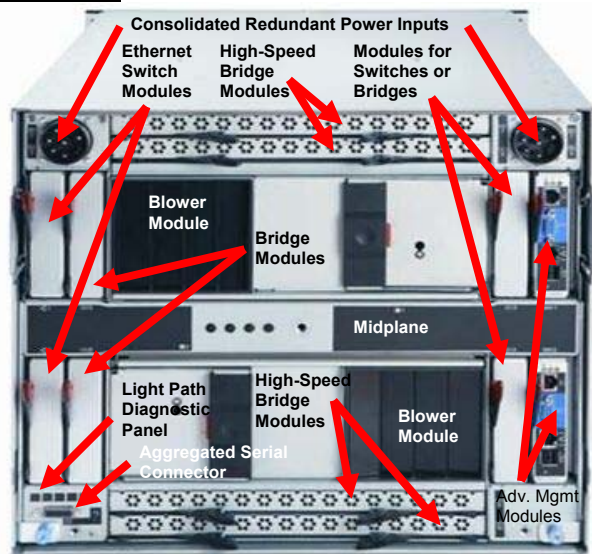| BladeCenter Chassis Specifications	
Machine type	**8677-3RU, 3RX**
Rack form factor (H x D)	7U x 28 inches
# of DVD/CD drives standard	1 **DVD-ROM** (in Media Tray)
# of diskette drives standard	1 (in Media Tray)
Internal tape drives supported	None (SAN-attached)
External disk drive support	NAS/SAN-attach
# of video ports	1 (on Management Module)
# of blade slots	**14 x 30mm** or **7 x 60mm** (or combinations thereof)
# of switch module slots	4 hot-swap/redundant
Switch modules standard	None (in standard chassis offerings)
# of management modules (std / max)	1 / 2 hot-swap/redundant
Systems management controller	Management Module
# of RS485 ports	1 (on Management Module)
# of serial breakout ports	None (Serial over LAN only)
# of parallel ports	None
# of mouse ports	1 (on Management Module)
# of keyboard ports	1 (on Management Module)
# of USB ports	2 (1 on Media Tray, one on Management Module)
Light path diagnostic panel	Yes (in Media Tray)
Predictive Failure Analysis support	Blades, switch modules, management modules, power modules, blower modules, media tray, HDDs
Power supply size standard	2000W **AC**
# of power supplies standard / maximum	2 / 4 hot-swap/redundant

16.

Please see the Legal Information section for important notices and information.

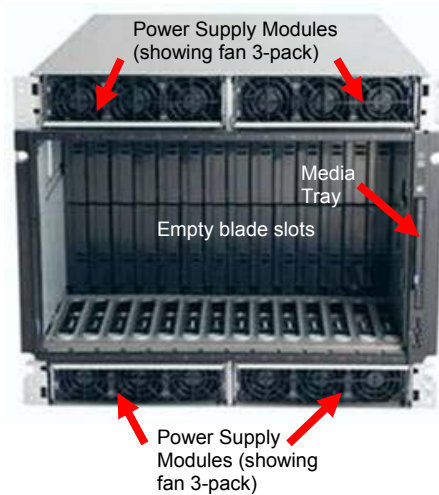| BladeCenter Chassis Specifications	
# of blower modules standard / max.	2 / 4 hot-swap/redundant
Blower module airflow (front-to-back)	250 CFM at 25ºC; 455 CFM at 32ºC or greater
Dimensions (HWD) / weight	12.0" (305mm) **H** 17.5" (4429mm) **W** 28.0" (711mm) **D** 99-240 **lbs** 44.9-108.9 **kg**
Length of limited warranty	3 years (parts and labor) onsite

BladeCenter H Chassis Images
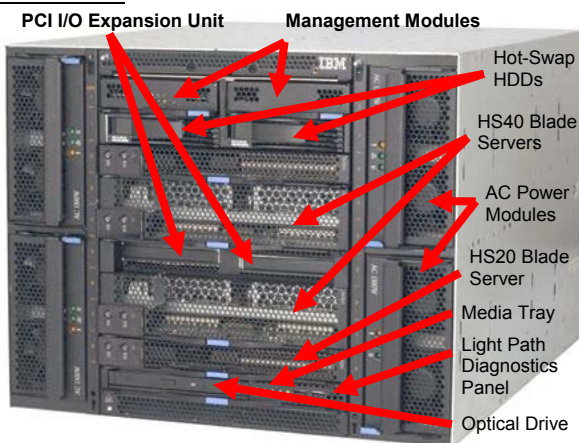
Front View

Rear View

Interior View

Please see the Legal Information section for important notices and information.
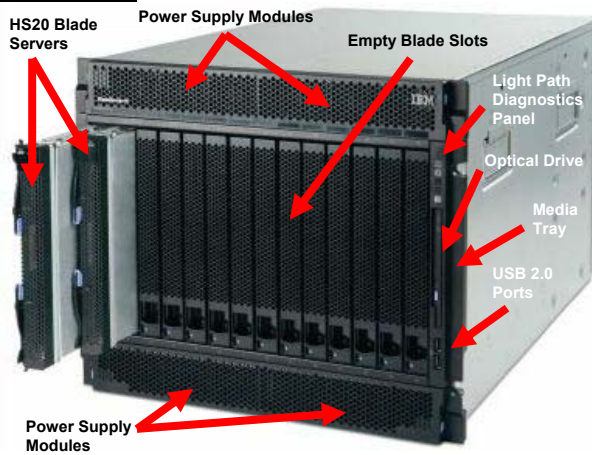
BladeCenter H Chassis Specifications	
Machine type	**8852-4XU**
Rack form factor (H x D)	9U x 28 inches
# of DVD-ROM drives standard	1 (in Media Tray)
# of diskette drives standard	None
Internal tape drives supported	None (SAN-attached)
External disk drive support	NAS/SAN-attach
# of video ports	1 (on Management Module)
# of blade slots	**14 x 30mm** or **7 x 60mm** (or combinations thereof)
Total # of switch/bridge module slots	10 (includes the following)
# of dedicated Gb Ethernet switch module slots	2 hot-swap/redundant
# of dedicated bridge module slots	2 hot-swap/redundant
# of switch / bridge module slots	2 hot-swap/redundant (either 2 switch or 2 bridge modules)
# of dedicated high-speed bridge module slots	4 hot-swap/redundant
# of communication / I/O modules standard	None
# of management modules (std / max)	1 / 2 hot-swap/redundant
Systems management controller	Advanced Management Module
# of RS485 ports	1 (on Management Module)
# of serial breakout ports	1
# of parallel ports	None (USB-based)
# of mouse ports	None (USB-based)
# of keyboard ports	None (USB-based)
# of USB 2.0 ports	2 (on Management Module)
Light path diagnostic panel	2 (one on Media Tray, one on Management Module)
Predictive Failure Analysis support	Blades, switch modules, bridge modules, management modules, power modules, blower modules, media tray, HDDs
Power supply size standard	2900W AC
# of power supplies standard / maximum	2 / 4 hot-swap/redundant
# of blower modules standard / max.	2 / 2 hot-swap/redundant
Blower module airflow (front-to-back)	430 CFM at 25ºC; 800 CFM at 32ºC or greater
Dimensions (HWD) / weight	15.75" (400mm) **H** 17.4" (442mm) **W** 28.0" (711mm) **D** — 41-159 **kg** 90-350 lbs
Length of limited warranty	3 years (parts and labor) onsite

BladeCenter T Chassis Images

Front View

PCI I/O Expansion Unit   Management Modules

Hot-Swap HDDs

HS40 Blade Servers

AC Power Modules

HS20 Blade Server

Media Tray

Light Path Diagnostics Panel

Optical Drive

Rear View

KVM Ports   Ethernet Switch Modules

Blower Module

Blower Module

Management Ports

Blower Module

Redundant Power Plugs

Blower Module

Ethernet Switch Modules
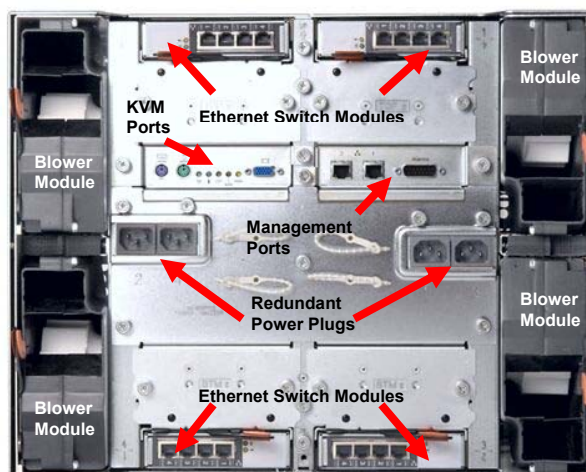
BladeCenter T Chassis Specifications	
Machine type	**8720-2RX (DC)** **8730-2RX (AC)**
NEBS/ETSI-compliance	Yes (NEBS Level 3)
Rack form factor (H x D)	8U x 20 inches
# of DVD/CD Combo drives standard	1 (in Media Tray)
# of diskette drives standard	1 (in Media Tray)
Internal tape drives supported	None (SAN-attached)
External disk drive support	NAS/SAN-attach
# of video ports	1 (in Media Tray)
# of blade slots	**8 x 30mm** (low-voltage), **7 x 30mm** (non-low-voltage), **4 x 60mm**, or combinations thereof
# of switch module slots	4
# of management modules std / max	1 / 2
Systems management controller	BladeCenter T Advanced Management Module
# of RS485 ports	None
# of serial breakout ports	None
# of parallel ports	None (USB-based)
# of mouse ports	None (USB-based)
# of keyboard ports	None (USB-based)
# of USB 2.0 ports	2 (in Media Tray)
Light path diagnostic panel	Yes (in Media Tray)
Predictive Failure Analysis support	Blades, switch modules, management modules, power modules, blower modules, media tray
Power supply size standard	1300W **DC** (8720-2RX), -38V to -75V (-48V nominal);

19.

Please see the Legal Information section for important notices and information.

BladeCenter T Chassis Specifications		
	1300W **AC** (8730-2RX)	
# of power supplies standard / maximum	2 / 4	
# of blower modules standard / max.	4 / 4	
Blower module airflow (front-to-back)	330 CFM (per blower)	
Dimensions (HWD) / weight	13.75" (349mm) **H** 17.4" (442mm) **W** 20.0" (508mm) **D**	98-192 **lb** 44.5-87.4 **kg**
Extended/Long Life Support	**3 yrs** production from date of GA, **5 yrs** post production support (standard)	
Length of limited warranty	3 years (parts and labor) onsite	

The Bottom Line

The HS20 offers maximum bang for the buck by incorporating outstanding features in a tiny package:

Performance:

- **Fast processor clock rates** — Up to two **3.8GHz single-core** 64-bit (EM64T) **Xeon** processors
- **Low-voltage processor clock rates** — Up to two **low-voltage 2.8GHz single-core** 64-bit (EM64T) **Xeon** processors
- **Dual-core processors** — Up to two **2.0GHz dual-core** 64-bit (EM64T) **Xeon** processors
- **Fast front-side bus** — **800MHz** FSB
- **Large cache** — **1MB** or **2MB** of L2 processor cache
- **Fast memory** — **400MHz PC2-3200 DDR II** ECC memory
- **Fast disk technology** — Internal **Ultra320 SCSI** or **SAS** storage (**2 HDDs**), with the option for two additional SCSI HDDs in an adjacent storage blade (**4 HDDs**) total
- **Fast I/O** —**InfiniBand, Myrinet** or **Fibre Channel** connectivity via an optional expansion card
- **Fast communications** — Integrated **dual Gigabit** Ethernet, with the option for two additional Gigabit Ethernet ports via an expansion card (total **4** per blade)

Flexibility:

- **Large memory capacity** — **16GB** using **four 4GB DIMMs**
- **High-capacity disk storage** — Up to **146.8GB** of internal **Ultra320 SCSI** storage, with the option for two additional SCSI HDDs in an adjacent storage blade (**4 HDDs, 746.8GB** total)
- **Integrated RAID** — **RAID-1** mirrored arrays standard; enhanced **RAID-1E** arrays optional (using a SCSI storage expansion unit)
- **One** *available* **adapter slot** standard —
  - ❑ One **legacy** slot for standard HS20 adapters, *or*
  - ❑ One **SPF** slot for new HS20 adapters
- Additional *optional* adapter slots —
  - ❑ *Two* **SFF** slots provided by the **SCSI Storage Expansion Unit**[6]
  - ❑ *Two* standard **PCI-X** slots provided by the **PCI I/O Expansion Unit II**[7]

Manageability and Availability:

- IBM **Director** systems management software
- Integrated **mini-Baseboard Management Controller** (mBMC):
  - ❑ **IPMI 1.5** compliance, including highly secure remote power control
- Interface to one or two **Management Modules** or **Advanced Management Modules** in the BladeCenter chassis for advanced systems management capability
  - ❑ Supports **LDAP** and **SSL** industry standards

---

[6] The SCSI Storage Expansion Unit requires an adjacent blade slot.

[7] The PCI/I/O Expansion Unit requires an adjacent blade slot.

- ❑ **Text console redirection** systems management
- ❑ **Serial over LAN**
- Numerous **hot-swap/redundant capabilities** provided via the chassis

Hot-swap SCSI storage via a SCSI Storage Expansion Unit

For More Information

IBM BladeCenter Servers and Options	**ibm.com**/systems/bladecenter
Competitive Sales Tool	w3.**ibm.com**/sales/systems/ibmsm.nsf/MainFrameset?OpenForm&cdoc=xscst
Customer References	w3.ncs.**ibm.com**/materials
Electronic Service Agent	**ibm.com**/support/electronic
Rack Configurator	**ibm.com**/pc/us/eserver/xseries/library/configtools.html
ServerProven Program	**ibm.com**/pc/us/compat
Technical Support	**ibm.com**/server/support
Other Technical Support Resources	**ibm.com**/servers/eserver/techsupport.html
Configuration and Options Guide	**ibm.com**/servers/eserver/xseries/cog

Legal Information

Warranty Information: For a copy of applicable product warranties, write to: Warranty Information, P.O. Box 12195, RTP, NC 27709, Attn: Dept. JDJA/B203. IBM makes no representation or warranty regarding third-party products or services including those designated as ServerProven or ClusterProven. Telephone support may be subject to additional charges. For onsite labor, IBM will attempt to diagnose and resolve the problem remotely before sending a technician.

IBM, the IBM logo, the e-business logo, AIX, BladeCenter, Calibrated Vectored Cooling, IBM System Storage, Predictive Failure Analysis, ServeRAID, System x and xSeries are trademarks of IBM Corporation in the United States and/or other countries. For a list of additional IBM trademarks, please see http://**ibm.com**/legal/copytrade.shtml.

AMD, the AMD arrow logo, AMD Opteron and combinations thereof are trademarks of Advanced Micro Devices, Inc.

InfiniBand is a trademark of the InfiniBand Trade Association.

Intel and Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds.

Microsoft, Windows and the Windows logo are trademarks or registered trademarks of Microsoft Corporation.

Other company, product and service names may be trademarks or service marks of others.

IBM reserves the right to change specifications or other product information without notice. References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates. IBM PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions; therefore, this statement may not apply to you.

This publication may contain links to third party sites that are not under the control of or maintained by IBM. Access to any such third party site is at the user's own risk and IBM is not responsible for the accuracy or reliability of any information, data, opinions, advice or statements made on these sites. IBM provides these links merely as a convenience and the inclusion of such links does not imply an endorsement.

Information in this presentation concerning non-IBM products was obtained from the suppliers of these products, published announcement material or other publicly available sources. IBM has not tested these products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Some machines are designed with a power management capability to provide customers with the maximum uptime possible for their systems. In extended thermal conditions, rather than shutdown completely, or fail, these machines automatically reduce the processor frequency to maintain acceptable thermal levels.

MB, GB and TB = 1,000,000, 1,000,000,000 and 1,000,000,000,000 bytes, respectively, when referring to storage capacity. Accessible capacity is less; up to 3GB is used in service partition. Actual storage capacity will vary based upon many factors and may be less than stated.

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will depend on considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

Maximum internal hard disk and memory capacities may require the replacement of any standard hard drives and/or memory and the population of all hard disk bays and memory slots with the largest currently supported drives available. When referring to variable speed CD-ROMs, CD-Rs, CD-RWs and DVDs, actual playback speed will vary and is often less than the maximum possible.

Please see the Legal Information section for important notices and information.